Folded Catadioptric Cameras*

Shree K. Nayar
Department of Computer Science
Columbia University, New York
nayar@cs.columbia.edu

Venkata Peri
CycloVision Technologies
295 Madison Avenue, New York
peri@cyclovision.com

Abstract

A framework is developed for the design and analysis of single-viewpoint catadioptric cameras that use two or more mirrors. The use of multiple mirrors permits folding of the optics which leads to more compact camera designs than ones that use a single mirror. A dictionary of camera designs that use two conic mirrors is presented. We show that any folded system that uses conic mirrors has a geometrically equivalent system that uses a single conic mirror. This result makes it easy to determine the scene-to-image mapping of a conic folded system. In addition, we discuss the optical benefits of using folded systems. As an example, we choose a camera design from our dictionary and optimize its parameters via optical simulations. This design is used to construct a compact video camera that provides a hemispherical field of view.

1 Introduction

Catadioptric cameras use a combination of mirrors and lenses to image the scene of interest. Of particular interest in computational vision are wide-angle cameras that satisfy the single viewpoint constraint; if a catadioptric system is capable of viewing the world from a single point in space, the captured image can be mapped to distortion-free images. Since such mapped images adhere to perspective projection, a variety of existing results in vision can be directly applied. Surveys of existing single-mirror catadioptric systems have been presented in [Nalwa, 1996] and [Nayar, 1997]. The complete class of single-mirror, single-lens imaging systems that satisfy the single viewpoint constraint have been analyzed in [Baker and Nayar, 1998].

A major issue with catadioptric imaging systems is that they tend to be physically large when compared with conventional ones. This is due to the fact that the capture of a wide unobstructed field of view requires the lens and the mirror to be adequately separated from each other. To work around this problem, the well-known method of optical folding is used. A simple example is the use of a planar mirror to fold the optical path between a curved mirror and an imaging lens. The folding can be in any direction; a 90° fold may help conceal some of the optical elements in an outdoor application and a 180° fold

reduces the size of the entire system. Folding by means of a curved mirror can result in greater size reduction. More importantly, curved folding mirrors can serve to reduce undesirable optical effects such as field curvature.

In the context of wide-angle imaging, a few folded systems have been implemented in the past. Here, we will focus on coaxial systems where the axes of the all the optical components are coincident. Buchele and Buchele [Buchele and Buchele, 1953] designed a single optical unit (a refractive solid) with a concave spherical mirror and a planar mirror attached to (or coated on) the solid. This idea was extended by Greguss [Greguss, 1986] who used a similar refractive solid with convex and concave aspherical mirrors. Powell [Powell, 1995] further improved the design by using a different shape for the refractive solid and convex and concave conics for the reflectors. Rees [Rees, 1970] has implemented a system that includes a convex hyperbolic primary mirror and a convex spherical secondary mirror. Rosendahl and Dykes [Rosendahl and Dykes, 1983] described an implementation that uses convex and concave hyperbolic mirrors and accompanying imaging optics for correction of field curvature. Charles [Charles, 1990] proposed a simple design in which the primary mirror is curved and the secondary one is planar. Several variants of the above designs have surfaced in the last few years which we will not review due to lack of space (see [Nayar and Peri, 1999] for recent implementations).

The previous work described above has not paid much attention to the single viewpoint constraint; the main objective has been to develop systems that produce high quality images of large fields of view. In this paper, we first look at the general problem of designing folded catadioptric cameras that have a single viewpoint. Geometric tools used in telescope design [Manly, 1991] and microwave optics [Cornbleet, 1984] are invoked in the context of wide-angle imaging. This leads to a general framework for designing multiple-mirror systems with single viewpoints. However, the mirror shapes are shown to be arbitrarily complex. Such mirrors make it difficult for the designer to minimize optical aberrations over the entire field of view. Hence, we restrict ourselves to designs that use conic mirrors whose optical manifestations are better understood and easier to correct. A complete dictionary of conic systems is presented within which some of the existing designs lie. In addition, we show that any folded system that uses conics can be geometrically represented

---

by an equivalent system that uses a single conic. This result makes is easy to determine the scene-to-image mapping of any folded system that uses conic mirrors.

Finally, as an example, we choose a specific design from our dictionary and implement a folded catadioptric video camera that is 9 cm tall, 5 cm wide and has a hemispherical field of view. The performance of the camera is described in terms of its spatially varying point blur function and enclosed energy plots. Perspective and panoramic video are shown that are computed from the hemispherical video using software.

2 Background: Single Mirror Systems

These cameras use a single mirror and a single lens to capture a large field of view that is observed from a fixed viewpoint. In [Baker and Nayar, 1998] the general problem of deriving mirror shapes that satisfy the fixed viewpoint constraint was studied. If $z(r)$ is the profile of the mirror shape, the complete class of solutions is given by

$$\left(z - \frac{c}{2}\right)^2 + r^2\left(1 - \frac{t}{2}\right) = \frac{c^2}{4}\left(\frac{t-2}{t}\right) ,$$

$$\left(z - \frac{c}{2}\right)^2 + r^2\left(1 + \frac{c^2}{2t}\right) = \left(\frac{2t + c^2}{4}\right) , \qquad (1)$$

where, c is the distance between the desired viewpoint and the entrance pupil of the imaging lens, and t is a constant of integration. This solution reveals that, to ensure a fixed viewpoint, the mirror must be a plane, ellipsoid, hyperboloid, or paraboloid (see [Baker and Nayar, 1998]).

3 Geometry of Folded Systems

As stated earlier, optical folding allows us to develop catadioptric cameras with significantly better packaging and optical performance.

3.1 The General Problem of Folding

The general problem of designing folded imaging systems can be stated as follows: Given a desired viewpoint location and a desired field of view, determine the mirrors (shapes, positions and orientations) that would reflect the scene through a single point, namely, the center of projection of the imaging lens. Though this problem has not be addressed in the context of wide-angle imaging, valuable theory has been developed for the construction of multiple-mirror telescopes and microwave devices [Cornbleet, 1984]. These are essentially imaging systems with very narrow fields of view (typically a couple of degrees). This theory is truly attractive in that it provides a suite of geometric tools for constructing folded systems (see [Cornbleet, 1984]). Here, we will outline the approach in the context of single-viewpoint, wide-angle systems.

Figure 1 shows a sketch of the problem. Let us assume that shape of the *primary mirror* is arbitrarily chosen and positioned with respect to the desired viewpoint O. Since the mirror has an arbitrary shape, the rays of light that travel from the scene towards the viewpoint O, after reflection by the mirror, do not necessarily converge at a

single point. Instead, they can be viewed as tangents to a surface that is called a *caustic*[1]. We would like to design a *secondary mirror* that would reflect these rays such that they intersect at a single point P, where the entrance pupil of the lens is located. For this, consider a string (dotted curve) with one end wound around and fixed to the caustic and the other end attached to the lens location P. Now, consider the point L that pulls on the string to keep it taut. As L moves along the string (while keeping it taut) in the direction shown in Figure 1, the string will wrap around the caustic and the locus of L is the required shape of the secondary mirror.
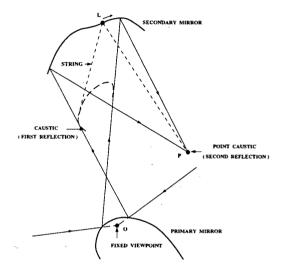

Figure 1: Geometrical construction of a wide-angle two-mirror imaging system. For any chosen primary mirror, a secondary mirror can be found that maps scene rays in the direction of a chosen viewpoint O to a chosen imaging pupil P.

It is worth noting that this elegant method for deriving mirror shapes from caustics can be applied repeatedly to design systems with more than two mirrors. It is a general technique for designing mirrors that transforms one caustic to another. In our case, the second caustic happens to be the point P. If the camera lens itself is not perspective but instead has a locus of viewpoints (yet another non-point caustic), it is possible to determine the secondary mirror that would map the first caustic to the second one, while ensuring that the complete system maintains a single viewpoint.

Clearly, the shape of the secondary mirror depends on the shape of the first caustic, which in turn depends on the shape of the primary mirror. Even for simple mirrors the caustics can have complex shapes such as nephroids, cardioids, cycloids, astroids, etc. For instance, in the case of collimated rays incident on a sphere, the caustic is a nephroid, which is rather complex [Cornbleet, 1984].

---

[1]Caustics have been used in vision for the recovery of specular shapes from highlights (see [Oren and Nayar, 1996]).

3.2 The Simpler World of Conics

As we have seen, a variety of exotic mirror pairs can be used to construct folded imaging systems with single viewpoints. However, complex mirror shapes tend to produce severe optical aberrations that cause image quality to vary dramatically over the field of view.

To keep geometrical and optical analysis simple we return to the conic mirrors given by equation (1). Note that each conic has well-defined foci that essentially serve as "point caustics" in relation to Figure 1. It is therefore easy to combine two (or more) conic mirrors to ensure a fixed viewpoint. To further simplify matters, we will restrict ourselves to coaxial imaging systems where the axes of the mirrors and the optical axis of the imaging lens coincide. A dictionary of the various configurations that result from using conic mirrors is show in Figure 2. Figure 2(a) shows a primary hyperboloidal mirror and a secondary planar mirror. Rays from the scene in the direction of near focus F_1 of the hyperboloidal mirror are reflected in the direction of its far foucs F_1'. The system is folded by placing the planar mirror between the near and far foci such that the far focus F_1' is reflected to the point P where the imaging lens is positioned, facing upward. The imaging lens and camera can therefore be placed inside the hyperboloidal mirror, further aiding compact packaging. Similarly, in Figure 2(b) the far focus of an ellipsoidal primary mirror is reflected by the planar mirror to P.

More sophisticated systems can be found in Figures 2(c)-(f) where the primary and secondary mirrors are hyperboloids and ellipsoids[2]. In each case, the near focus of the secondary mirror is made to coincide with the far focus of the primary mirror. The entrance pupil of the imaging system is then placed at the far focus of the secondary mirror. Figure 2(g) shows how a concave hyperboloid may be used. A few more systems that use concave hyperboloids and convex ellipsoids exist but are omitted for brevity.

Finally, in Figure 2(h) and (i) paraboloidal primary and secondary mirrors are used. In these cases, the primary mirror orthographically reflects all rays of light incident in the direction of its focus F_1. These rays are collected by a secondary paraboloid and reflected so as to converge at its focus F_2, where the lens is positioned. In effect, the secondary mirror and the imaging lens together serve as a telecentric imaging system as used in [Nayar, 1997].

3.3 Equivalent Single Mirror Systems

Here, we show that any folded system with two conic mirrors can be geometrically represented by an equivalent system with a single conic mirror, where the scene-to-image mapping of the original system is preserved by the equivalent one. It should be noted that geometrical equivalence does not imply optical equivalence. Even so, the geometrical equivalence is valuable in that it enables one to easily determine the relation between scene points

---
[2]This particular combination was pointed out by Sergey Trubko [Trubko, 1998] at CycloVision Technologies.

and image coordinates, which is needed to map images produced by a folded system to perspective or panoramic ones. Our equivalence proof here will be brief (see [Nayar and Peri, 1999] for details).

Figure 3 shows a sketch of a folded system with two conic mirrors. Since the system has axial symmetry, the equivalence need be established only for a one-dimensional cross-section. Let the primary mirror C_1 have conic constant k_1, radius of curvature R_1, and near and far foci F_1 and F_1'. The shape of the mirror is fully determined by its conic constant: $k_1 = 0$ gives a sphere, $0 > k_1 > -1$ yields an ellipsoid, $k_1 = -1$ gives a paraboloid and $k_1 < -1$ results in a hyperboloid. The distance between its foci is $2R_1\sqrt{-k_1}/(1 + k_1)$.

While the conics in equation (1) are defined with the near focus at the origin, we can move the origin to the apex to get the simpler form:

$$r^2 = -2R_1 z - (1 + k_1)z^2 . \qquad (2)$$

The same in polar coordinates, with the origin at the near focus F_1, is

$$\rho = \frac{R_1}{1 + \sqrt{-k_1}\cos\theta} , \qquad (3)$$

where

$$z = \rho\cos\theta - R_1/(1 + \sqrt{-k_1}) , \quad r = \rho\sin\theta . \qquad (4)$$

The scene ray L_1 in the direction of F_1 strikes the primary mirror C_1 at P_1. The slope of the mirror at P_1 is

$$m_1 = \frac{dr}{dz} = \frac{-R_1 - (1 + k_1)z}{r} . \qquad (5)$$

Using (4), we can substitute for z and r to get

$$m_1 = -\frac{\sqrt{-k_1} + \cos\theta}{\sin\theta} . \qquad (6)$$

From this expression for the slope and the specular reflection constraint (incidence angle equals reflection angle), we get the following simple relation between the angle θ of an incoming scene ray L_1 and the angle α of the reflected ray L_2:

$$\tan\alpha = \frac{(1 + k_1)\sin\theta}{2\sqrt{-k_1} + (1 - k_1)\cos\theta} . \qquad (7)$$

The above expression determines the *compression* of rays due to the primary mirror. The secondary mirror C_2 is also a conic (with constant k_2). Since its near focus F_2 coincides with the far focus F_1' of the primary mirror C_1, the rays reflected by C_1 are directed towards F_2. Hence, the above compression equation can be used to relate the angle α of an incoming ray L_2 to the angle β of the reflected ray L_3:

$$\tan\beta = \frac{(1 + k_2)\sin\alpha}{2\sqrt{-k_2} + (1 - k_2)\cos\alpha} . \qquad (8)$$

219

(a)

F1' •
PL — (2)
HYP
P
F1
(1)

(b)

F1' •
PL — (2)
ELL
F1
P
(1)

(c)

F2, F1'
HYP (2)
HYP
P, F2'
F1
(1)

(d)

ELL (2)
F2, F1'
HYP
P, F2'
F1
(1)

(e)

ELL (2)
F2, F1'
ELL
F1
P, F2'
(1)

(f)

F2, F1'
HYP (2)
ELL
F1
P, F2'
(1)

(g)

ELL (2)
F2, F1
HYP
P, F2'
F1'
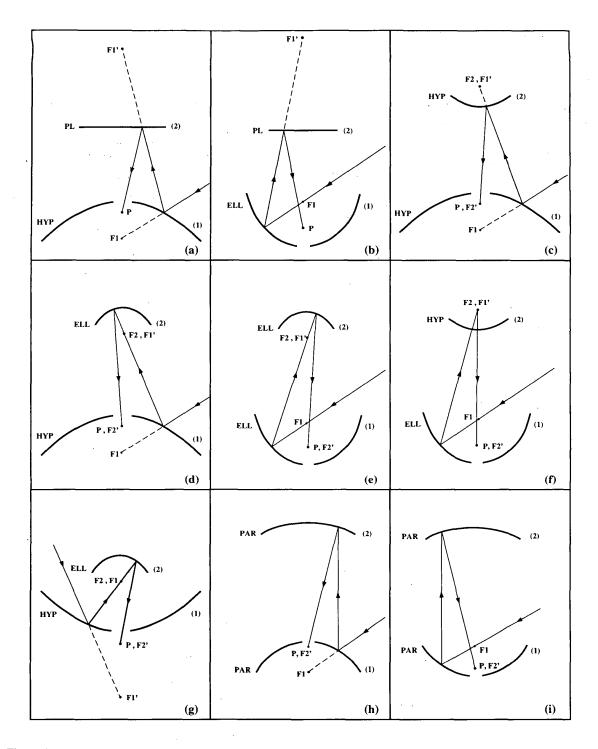
(h)

PAR (2)
PAR
P, F2'
F1
(1)

(i)

PAR (2)
PAR
F1
P, F2'
PAR (1)

Figure 2: A dictionary of two-mirror folded catadioptric camera designs that satisfy the single viewpoint assumption. In this dictionary only mirrors with conic cross-sections are used. Mirrors with the following shapes are used: planar (PL), hyperboloidal (HYP), ellipsoidal (ELL) and paraboloidal (PAR). The primary and secondary mirrors are denoted by (1) and (2), respectively. The near and far foci of the primary mirror are denoted by $F1$ and $F1'$, and those of the secondary mirror by $F2$ and $F2'$. The entrance pupil of the imaging lens is positioned at P.
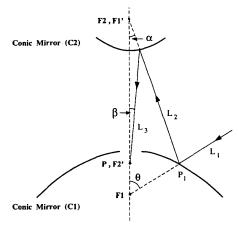
220

Figure 3: Any single-viewpoint folded system that uses two or more conic mirrors has an equivalent single mirror system with the same compression, which is the relation between the directions of scene points (θ) and their image coordinates (determined by β).

From equations (7) and (8) we get the compression of the complete folded system:

$$\tan\beta = \qquad (9)$$
$$\frac{(1 + k_1)(1 + k_2)\sin\theta}{\left[2(\sqrt{-k_1} + \sqrt{-k_2} - k_1\sqrt{-k_2} - \sqrt{-k_1 k_2})\right.}$$
$$\left. + (1 + k_1(k_2 - 1) + 4\sqrt{-k_1}\sqrt{-k_2} - k_2)\cos\theta\right]$$

If neither mirror is a paraboloid, i.e. $k_1 \neq -1$ and $k_2 \neq -1$, the above compression is the same as that produced by a single conic mirror with a conic constant of either k_e or $1/k_e$ where [3]

$$k_e = -\left(\frac{\sqrt{-k_1} + \sqrt{-k_2}}{1 + \sqrt{-k_1}\sqrt{-k_2}}\right)^2 . \qquad (10)$$

For each of the folded configurations shown in Figure 2(c)-(g), the equivalent conic is either a hyperboloid or an ellipsoid. The equivalent conic is a sphere for the special (but impractical) case of a folded system made of two concentric spheres.

The paraboloidal configurations ($k_1 = k_2 = -1$) in Figures 2(h) and 2(i) also have equivalent single-conic systems. In these cases, k_e is a function of the *parameters* h_1 and h_2 of the two paraboloids and can be shown to be [Nayar and Peri, 1999]

$$k_e = -\left(\frac{h_1 + h_2}{h_1 - h_2}\right)^2 . \qquad (11)$$

Here again, the equivalent conic can be an ellipsoid or a hyperboloid when $h_1 \neq h_2$. When the two paraboloids

[3] In equation (7) we see that for $k_1 \neq 0$ and $k_1 \neq -1$, $\tan\alpha|_{k_1,\theta} = -\tan\alpha|_{1/k_1,\theta}$. That is, the compression by an ellipsoid of conic constant k_1 is equivalent to the compression by a hyperboloid of conic constant $1/k_1$.

are identical, i.e. $h_1 = h_2$, no compression of the field of view is achieved and the equivalent conic is a sphere with the viewpoint at its center.

4 Optics of Folded Systems

The above designs only define the geometry of the sensor. That is, the entrance pupil of the imaging system is taken to be a pinhole and hence only the principal rays are considered. When a lens is used to gather more light, each principal ray is accompanied by a bundle of surrounding rays and a variety of optical aberrations appear that make the design of a folded system challenging.

4.1 Pertinent Optical Effects

Here, we briefly describe some of the optical aberrations that are pertinent to us (see [Hecht and Zajac, 1974] for details).

Chromatic Aberration: The focal length of any lens will vary somewhat with the "color" of the incoming light. An imaging lens will have several individual elements and one of the design goals is to ensure that chromatic aberrations induced by individual elements at least partially compensate for each other.

Coma and Astigmatism: Both these aberrations are caused primarily due to the curvatures of the mirrors. The effect of coma is proportional to square of the aperture size, while astigmatism is linear in the aperture size. Both effects cause the best focused image of a scene point to not be a single point but rather a volume (of confusion). Our design goal is to maximize aperture size (minimize F-number) while ensuring that the blur function falls within a single detector (pixel) for all points in the field of view.

Field Curvature: Since at least one of our mirrors is curved, points at infinity end up being best focused not on a plane but rather a curved surface behind the imaging lens. This curved surface is also called the Petzval surface [Hecht and Zajac, 1974]. Since the CCD imagers we have at our disposal are planar, the best image quality is achieved where the curved image and the planar detector intersect. In compact systems (small mirrors with high curvatures) field curvature tends to dominate over all other aberrations. In a single mirror system, the image surface is curved in the same direction as the mirror itself. Hence, in a two-mirror system it is to our advantage to use a convex and a concave mirror so that the field curvatures introduced by the two mirrors serve to compensate for each other.

4.2 Design Parameters

Thus, the design of a catadioptric system requires the selection of optical parameters that minimize a variety of complex aberrations. Before we describe how the optimization is performed, let us summarize the parameters at hand. Since the total number of parameters are generally very large, it helps to fix some of them prior to system optimization.

CCD Size: A few different CCD formats are commercially available (1 inch, 1/2 inch, 1/3 inch, 1/4 inch, etc.). If the number of pixels in each CCD is more or less the same, the pixel size reduces with CCD size. Typically,

the choice of the CCD format is based on the packaging and resolution requirements of the application.

Imaging Lens: The parameters of the imaging lens are characterized by its focal length, field of view, number of elements and its speed (aperture size). While the number of elements and their basic shapes (convex, concave, meniscus, etc.) may be selected up-front by the designer, the curvatures and diameters of the lenses may be treated as free parameters. Once the optimization is done, one tries to match the resulting parameters with those of commercially available lenses.

Mirrors: As we have seen in section 3, a large number mirror shapes are feasible from the perspective of geometry. Based on the size and field of view requirements, as well as a good deal of intuition, one must select the general shapes of the mirrors to be used. Further, since we know *apriori* that the use of a convex and a concave mirror helps to reduce field curvature, such a choice can be made up-front. Once the basic shapes have been chosen, the exact shape parameters (conic constants, for instance) can be treated as free parameters to be optimized.

Distances: We know that to achieve a single viewpoint, the far focus of one mirror must coincide with the near focus of the other. In addition, fairly tight bounds on the distances between the individual optical components can be given based on the sensor size requirements imposed by the application. The exact distances can then be treated as free parameters in the optimization process.

4.3 System Optimization

In our work, the free parameters are computed using the Zemax software package from Focus Software Incorporated. The package performs iterative numerical optimization using fast ray-tracing. During each iteration, images of point sources in the scene are generated. An objective function is formulated to yield a minimum when the ray-traced point spread functions are most compact.

5 An Example Implementation

As an example, we will describe a folded panoramic video camera we have implemented. The camera uses the layout shown in Figure 2(h), wherein two parabolic mirrors are used. Note that the secondary mirror has a significantly longer focal length than the first one. This is because the two mirrors must be adequately separated to avoid a large blindspot due to obstruction by the secondary mirror. Prior to optimization, it was decided that the complete sensor must lie within a cylinder that is 90 mm tall and 50 mm in diameter. The desired field of view was set to a hemisphere and the maximum allowable blindspot to 22 degrees when measured from the optical axis. It was also decided that a 1/3 inch CCD camera would be used. Given these constraints, the secondary mirror ends up being a small (shallow) section of a paraboloid, which is well-approximated by a spherical mirror. Using the above numbers as upper bounds, the parameters of the entire system were optimized.

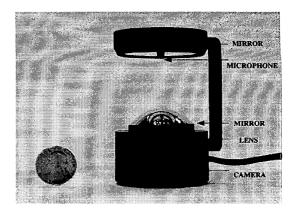Figure 4 shows the resulting device. The primary

Figure 4: A folded catadioptric camera with a hemispherical field of view. The device is 90 mm tall and 50 mm wide. It includes folded optics, a video camera and a microphone.

parabolic mirror has a focal length of 10 mm and a diameter of 40 mm. The secondary spherical mirror has a radius of curvature of 110 mm. The video camera used is a Computar EMH200 board camera with 550 horizontal lines of resolution, and the imaging lens has a focal length of 6mm. Finally, a microphone is attached to the center of the secondary mirror, pointing towards the primary mirror. This effectively maps the narrow response cone of the microphone to a panoramic one.
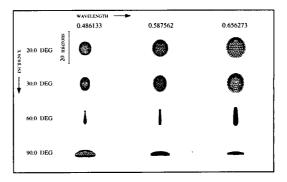

Figure 5: Spot diagrams for various wavelengths (columns) and angles of incidence (rows), computed using the optimized optical design for the camera shown in Figure 4.

Figure 5 shows the matrix spot diagram for the above design. Each spot can be viewed as the point blur function for a specific wavelength of light (columns) and a specific angle of incidence (rows). The scale bar shown beside the top-left spot is 20 microns long. As seen, the spots vary in shape quite a bit. This is due to aberrations caused by coma, astigmatism, field curvature and chromatic aberration. The goal of the optimization was to ensure that all the spots (across the different wavelengths and angles of incidence) are kept as compact as possible. Figure 6 shows the energy plots for the different angles of incidence. As the dotted lines indicate, for all angles of incidence, about 70% of the total energy in the point spread function lies within a circle of radius 4 microns. In our
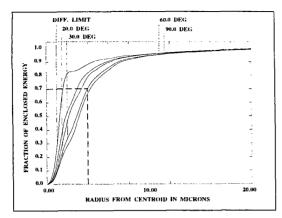
Figure 6: Encircled energy plots for different angles of incidence, for the camera shown in Figure 4 (see text for details).

case, the pixel size on the CCD chip is approximately 6.4 x 7.4 microns. Hence, the above system produces an almost ideal digital image.

Figure 7(a) shows an image captured using the sensor. As can be seen, despite all the complex optical aberrations at work, the sensor produces a clear image that has a very large depth of field for all angles of incidence. Figures 7(b) and (c) show perspective and panoramic video streams that are computed from the hemispherical video.

Acknowledgments
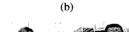
References

[Baker and Nayar, 1998] S. Baker and S. K. Nayar. Catadioptric Image Formation. *Proc. of International Conference on Computer Vision*, January 1998.

[Born and Wolf, 1965] M. Born and E. Wolf. *Principles of Optics*. London:Permagon, 1965.

[Buchele and Buchele, 1953] D. R. Buchele and W. M. Buchele. Unitary Catadioptric Objective Lens Systems. *United States Patent*, (2,638,033), May 1953.

[Charles, 1990] J. R. Charles. Portable All-Sky Reflector with "Invisible" Axial Camera Support. *United States Patent*, (D312,263), November 1990.

[Cornbleet, 1984] S. Cornbleet. *Microwave and Optical Ray Geometry*. John Wiley and Sons, 1984.

[Greguss, 1986] P. Greguss. Panoramic Imaging Block for Three-Dimensional Space. *United States Patent*, (4,566,763), January 1986.

[Hecht and Zajac, 1974] E. Hecht and A. Zajac. *Optics*. Addison Wesley, Reading, Massachusetts, 1974.

[Manly, 1991] P. L. Manly. *Unusual Telescopes*. Cambridge University Press, 1991.

[Nalwa, 1996] V. Nalwa. A True Omnidirectional Viewer. Technical report, Bell Laboratories, Holmdel, NJ 07733, U.S.A., February 1996.

(a)

(b)

(c)

Figure 7: (a) Hemispherical video produced by the catadioptric camera shown in Figure 4. Software is used to map the hemispherical video to (b) perspective and (c) panoramic video streams. The jaggy artifacts are due to the low resolution (640x480) of the original video.

[Nayar and Peri, 1999] S. K. Nayar and V. Peri. Folded catadioptric imaging systems. *Technical Report, Department of Computer Science, Columbia University*, May 1999.

[Nayar, 1997] S. K. Nayar. Catadioptric Omnidirectional Camera. *Proc. of IEEE Conf. on Computer Vision and Pattern Recognition*, June 1997.

[Oren and Nayar, 1996] M. Oren and S. K. Nayar. A Theory of Specular Surface Geometry. *International Journal of Computer Vision*, 24(2):105–124, 1996.

[Powell, 1995] I. Powell. Panoramic Lens. *United States Patent*, (5,473,474), December 1995.

[Rees, 1970] D. W. Rees. Panoramic Television Viewing System. *United States Patent*, (3,505,465), April 1970.

[Rosendahl and Dykes, 1983] G. R. Rosendahl and W. V. Dykes. Lens Systems for Panoramic Imagery. *United States Patent*, (4,395,093), July 1983.

[Trubko, 1998] S. Trubko. *Personal Communication*. August 1998.